

The Case for Case-Based Transfer Learning

Matthew Klenk, David W. Aha, and Matt Molineaux

■ *Case-based reasoning (CBR) is a problem-solving process in which a new problem is solved by retrieving a similar situation and reusing its solution. Transfer learning occurs when, after gaining experience from learning how to solve source problems, the same learner exploits this experience to improve performance and learning on target problems. In transfer learning, the differences between the source and target problems characterize the transfer distance. CBR can support transfer learning methods in multiple ways. We illustrate how CBR and transfer learning interact and characterize three approaches for using CBR in transfer learning: (1) as a transfer learning method, (2) for problem learning, and (3) to transfer knowledge between sets of problems. We describe examples of these approaches from our own and related work and discuss applicable transfer distances for each. We close with conclusions and directions for future research applying CBR to transfer learning.*

Observations of human reasoning motivate AI research on transfer learning (TL) and case-based reasoning (CBR). Our ability to transfer knowledge and expertise from understood domains to novel ones has been thoroughly documented in psychology and education (for example, Thorndike and Woodworth 1901; Perkins and Salomon 1994; Bransford, Brown, and Cocking 2000), among other disciplines. Transfer learning uses knowledge learned from solving tasks from a source domain to enhance an agent's ability to learn to solve tasks from a target domain. The differences between the source and target problems characterize the *transfer distance*. Case-based reasoning transfers problem-solving knowledge from specific examples or episodes, called *cases*, to new problems.

While researchers typically work within each of these fields independently, the purpose of this article is to summarize how case-based reasoning can be applied to transfer learning.¹ Our analysis reveals three approaches for applying CBR to transfer learning: (1) CBR as a transfer learning method, (2) CBR for problem solving, and (3) CBR to transfer knowledge between the domains. These correspond to using CBR for solving entire transfer tasks or acting as a component within a transfer learning system. The transfer distance provides a new metric for CBR researchers to assess the robustness of their systems. Furthermore, each CBR approach has implications for the importance and interpretations of different transfer learning metrics.

We begin with an overview of transfer learning, case-based reasoning, and the three approaches for applying CBR to transfer learning. We provide examples of each from our research on physics problem solving and controlling a player in a football simulation and also describe related work. We close with a discussion concerning the applicability of each approach over different transfer distances, how transfer learning assists CBR research, and some directions for future research.

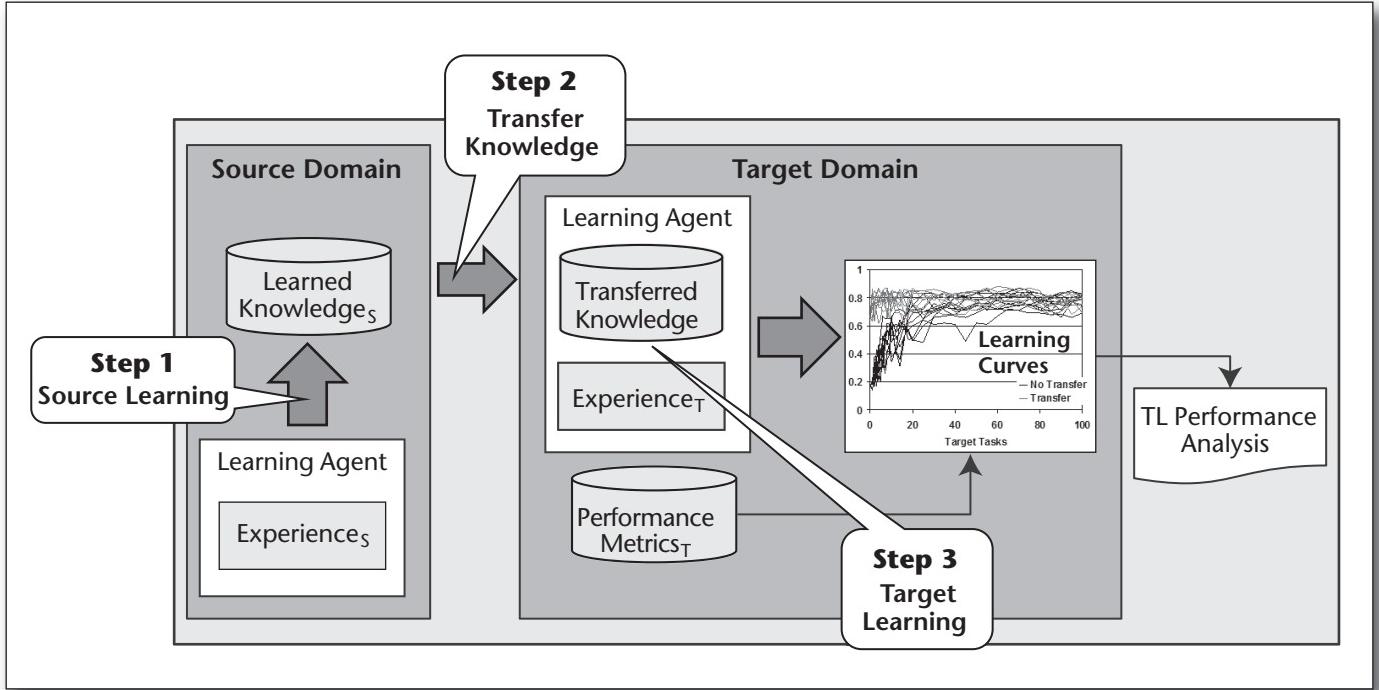


Figure 1. Transfer Learning Framework and Evaluation.

Combining Two Methods for AI Research

In this section, we first describe the transfer learning and case-based reasoning frameworks. Using these frameworks, we define the three approaches for applying CBR in transfer learning systems at the end of this section.

Transfer Learning

Machine learning has traditionally focused on learning from a blank slate in isolated tasks. This differs substantially with the way humans learn; people frequently leverage experience gained from learning one task to improve their performance on different, novel tasks. Transfer learning is the process of recognizing and applying knowledge and skills learned from one or more previous (source) problems to more efficiently or effectively learn to solve novel (target) problems. Methods for transfer learning hold the promise of being exceedingly useful; they could dramatically decrease the amount of training required by successfully employing knowledge obtained from different, but related, problems. This promise motivates the development of computational models for transfer learning, which has been the focus of workshops at NIPS-05, ICML-06, AAAI-08, and NIPS-09, in addition to a large amount of work summarized in recent surveys (Taylor and Stone 2009; Torrey and Shavlik 2009; Pan and Yang 2010).

Figure 1 summarizes a conceptual model of

transfer learning and its empirical evaluation, where the domain representations, tasks, performance metrics, and environments may all differ between the source and target problems. Transfer learning involves three steps: (1) learning in the source, (2) transferring the learned knowledge from source to the target, and (3) learning in the target.

A transfer learning evaluation compares an agent's performance on a task defined in a target domain after learning on a task or tasks from the source (the transfer condition) to the agent's performance without any source experience (the non-transfer condition). Common performance measures include initial advantage, learning rate, and asymptotic advantage. *Initial advantage (or jump start)* is the initial increase in an agent's performance resulting from transfer. *Learning rate* is a decrease in the time required to reach a particular performance level, particularly asymptotic performance. This is usually measured using k -step regret (Kaelbling, Littman, and Moore 1996). *Asymptotic advantage* is that the agent's final performance may be improved through transfer.

Differences between the source and target problems may be categorized by their transfer distance. Defining common transfer distances across a wide range of tasks, for example, physics problem solving and strategy games, is quite challenging. Typically, individual researchers define transfer distance for particular evaluations. For example, Pan and Yang (2010) proposed a new categorization based on two assumptions made in traditional

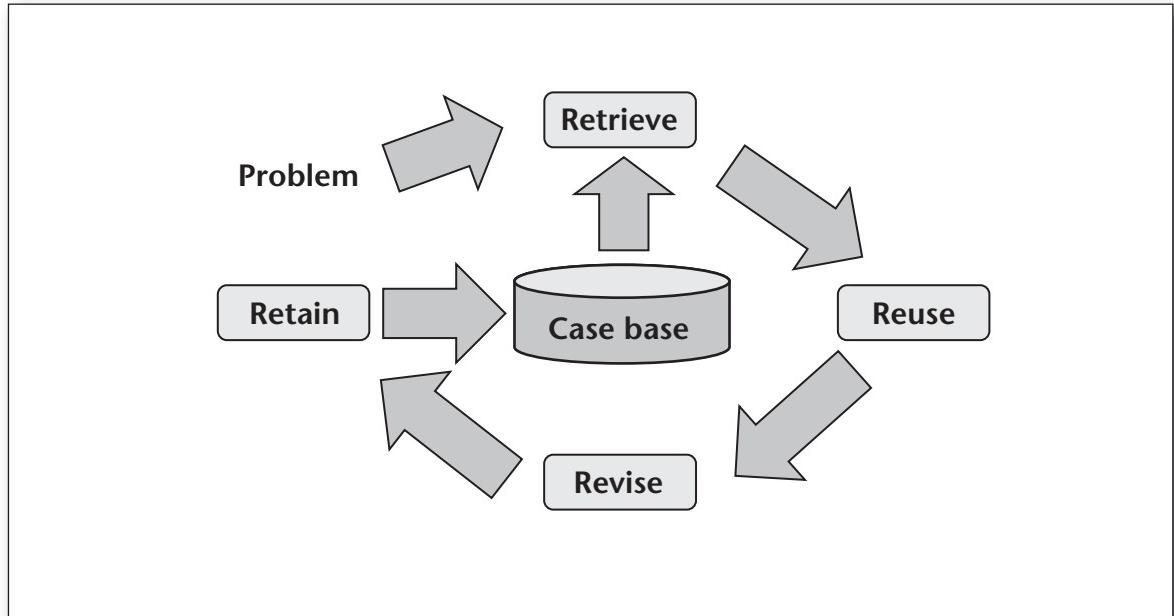


Figure 2. The Case-Based Reasoning Cycle.

Adapted from López de Mántaras et al. (2005).

machine-learning studies. In particular, traditional machine learning assumes that the task to be performed (that is, classes and class objective functions) and the domain of performance (that is, feature spaces and instance distributions) do not change among the source and target problems. In transfer learning, these assumptions are relaxed; the source and target problems may involve different tasks and domains. Refining definitions of transfer distance is essential to maturing the field of transfer learning, and we refer to this issue repeatedly in this article.

Case-Based Reasoning

Case-based reasoning is a problem-solving process in which inferences about a situation are drawn from individual instances called cases. While the roots of CBR lie in observations of human reasoning (Schank 1982; Kolodner 1993), this discipline is now aligned closely with computer science. CBR research today focuses on the study of algorithms, representations, and their applications for a large variety of analysis and synthesis tasks (Aamodt and Plaza 1994; López de Mántaras et al. 2005). For example, there have been significant CBR contributions on recommender systems, legal reasoning, textual reasoning, and planning tasks (Aha, Marling, and Watson 2005). Current research venues include a dedicated conference (ICCBR), several annual CBR-focused workshops, and other AI and cognitive science venues.²

The traditional case-based problem-solving cycle, shown in figure 2, includes four steps. Given

a new problem p to solve, this process begins by retrieving one or more similar cases (that is, problem-solution pairs) from its memory, or case base, and reusing their solutions. Solution reuse may require an adaptation of the retrieved cases' solutions for use in solving p . The proposed solution s may be revised in light of its performance. Finally, the pair (p, s) may be retained as a new case. In this manner, CBR systems learn by accumulating cases and refining models that support their four steps.

Theories of analogy also computationally explore reasoning from cases. The structure-mapping theory (Gentner 1983) defines analogy as an alignment process between two structured representations resulting in inferences about the target from the source. This alignment process may be used for retrieval and reuse within a CBR system. We introduce analogy, and its connection to CBR, here because it plays an important role in a number of the transfer learning methods discussed in this article, either within CBR systems or as basis for potential mappings between domains.

Approaches to Transfer Learning Using Case-Based Reasoning

Recall transfer learning involves three steps: learning in the source, transferring the learned knowledge from source to the target, and learning in the target. We categorize CBR approaches for transfer learning by which transfer learning steps are performed by the CBR cycle. We identify three types of approaches, shown in table 1, for applying CBR in transfer learning systems: (1) as a transfer learn-

	Source Learning	Knowledge Transfer	Target Learning	Examples
CBR as a transfer learning method	✓	✓	✓	Hinrichs and Forbus (2007), Sharma et al. (2007), Klenk and Forbus (2009)
CBR for problem learning	✓	✗	✓	Wu and Dietterich (2004), Shi et al. (2009), Aha et al. (2009)
CBR to transfer knowledge	✗	✓	✗	Lui and Stone (2006), Kuhlmann and Stone (2007), Shi et al. (2009), Konik et al. (2009), Hinrichs and Forbus (2011)

Table 1. Applications of CBR in Transfer Learning Systems.

ing method, (2) for problem learning, and (3) to transfer knowledge. After defining each approach, we present examples of them from our research and related work.

CBR as a Transfer Learning Method. CBR can be used directly as a transfer learning method. Recall that all CBR methods involve transferring knowledge from prior cases to new problems. In CBR as a transfer learning method, the CBR cycle accounts for all three steps of transfer learning. For example, the learned and transferred knowledge could be the case base after training on source problems. During learning on target problems, the same CBR cycle can be used to solve problems in the target, updating the same case base. Thus, the CBR system is unaware that it is being evaluated for transfer learning and makes no distinction between source and target cases. For these systems, transfer distance and initial advantage provide a useful metric for evaluating the retrieval and reuse mechanisms of the CBR system.

CBR for Problem Learning. CBR can be used for problem learning (source, target, or both). In this approach, the source and target cases are separated and treated distinctly by the system. To be a transfer learning system, the CBR system must be integrated with another component to perform knowledge transfer between the source and target problems. Transfer learning provides structure guiding the construction and evaluation of integrated CBR systems.

CBR to Transfer Knowledge. CBR methods can be used for transferring knowledge from source to target. These approaches either use a full CBR cycle to modify source instances for use in target learning by another algorithm, or they create analogical mappings between problem domains to support the transfer of abstract knowledge. These methods must be integrated with a learning mechanism to perform transfer learning.³ Thus, these approaches convert established learning methods, such as,

hierarchical skill learning (Nejati, Langley, and Könik 2006) or Q-learning (Watkins 1989), into transfer learning methods. These methods support transfer distances in which the problem and solution representations for the source and target include different relations, types, and quantities.

In the next three sections, we discuss each method in more detail along with an example from our research.

CBR as a Transfer Learning Method: AP Physics

To illustrate how a CBR method may be directly used as a transfer learning method, we describe Klenk and Forbus's approach for advanced-placement (AP) physics problem solving (2009a). Physics problem solving requires reasoning over a wide range of entities and scenarios. While the authors sidestep natural language understanding by using predicate-calculus representations, the translation process leaves the everyday concepts in place. That is, balls, buildings, astronauts, boxes, baseball bats, flying, falling, and pulling all appear in the formal problem descriptions. They used a subset of the ResearchCyc (research.cyc.com) ontology containing over 30,000 concepts. Understanding the relevant abstractions and assumptions for a physics problem stated as an everyday situation is a difficult problem, in part because modeling decisions are contextual. For example, a coin falling off a building can be considered to be a point mass. However, when modeling the exact same coin spinning on a table, it cannot be considered a point mass since its shape and size must be considered.

Solving Physics Problems Using Analogical Model Formulation

Given a scenario from a physics problem, an intelligent problem solver can use model formulation (Falkenhainer and Forbus 1991) to construct a sce-

nario model, which consists of the relevant abstractions, assumptions, and equations necessary for answering the question. An important contribution of the qualitative reasoning community has been formalizing this process (Rickel and Porter 1994). While successful in engineering applications, these approaches are limited in that they focus on abstract scenarios, require complete and correct domain theories, and ignore learning. To overcome these limitations, analogical model formulation, or AMF (Klenk and Forbus 2009a), builds scenario models of everyday situations based on experiences. Analogical model formulation incrementally learns by accumulating examples and making effective use of them, even when its knowledge is incomplete.

Analogical model formulation was evaluated as a transfer learning method for advanced-placement physics problem solving using 460 AP physics-style problems created by the Educational Testing Service (ETS) and Cycorp. These problems were generated using variations of four problem types typically found on the AP physics exam. The source consisted of 20 problems. Using these problems, ETS created sets of problems to exemplify six distinct transfer distances representing systematic differences between source and target problems: parameterization (changing the parameter values, but not the qualitative outcome), extrapolation (changing the parameters such that the qualitative outcome changes as well), restructuring (asking for a different parameter), extending (including distracting information), restyling (changing the types of everyday objects involved), and composing (requiring concepts from multiple source problems).

When learners study for the AP physics exam, one important way in which they learn is by solving problem sets. For feedback, they often get worked solutions—step-by-step explanations typically found in the back of textbooks. AMF performs CBR by employing worked solutions as cases. When presented with a new problem, AMF uses the many are called but few are chosen structure-mapping engine (MAC/FAC) (Forbus, Gentner, and Law 1995) to retrieve an analogous worked solution and SME (Falkenhainer, Forbus, and Gentner 1989) to reuse its modeling decisions to construct a scenario model of the new problem. MAC/FAC selects an analogous case from the case base in two stages: (1) a nonstructural match using feature vectors whose weights are proportional to the number of occurrences of each predicate in a representation and (2) a structural alignment to determine its relational similarity. In the first stage, the problem is compared to each case in the case base with the three most similar selected for structural discrimination. To determine the closest structural match, SME computes match scores by creating an ana-

logical mapping between each selected case and the problem, and the closest match is selected. Analogical mappings consist of correspondences between the entities and expressions of the worked solution and problem. From these correspondences, SME creates a set of candidate inferences, which are conjectures about the problem using expressions from the base, which, while unmapped in their entirety, have subcomponents that are included in the correspondences.

Analogical model formulation is implemented in the Companion cognitive architecture (Forbus, Klenk, and Hinrichs 2009), which is exploring the hypothesis that analogical processing is central to human reasoning and learning. After attempting a problem, the Companion is provided with its worked solution, which it retains for future problem solving by adding it to its case base.

Central to this approach is the use of candidate inferences generated by the analogical mapping to make modeling decisions. Consider the following source problem:

An astronaut on a planet with no atmosphere throws a baseball upward from near ground level with an initial speed of 4.0 meters per second. If the baseball rises to a maximum height of 5.0 meters, what is the acceleration due to gravity on this planet? (a) 0.8 m/s²; (b) 1.2m/s²; (c) 1.6m/s²; (d) 20m/s²

The worked solution to this problem includes seven steps instantiating the relevant equations, assuming parameter values and solving equations. Figure 3 includes a subset of the representation, simplified for presentation, of one step from the worked solution and a restyling problem in which the baseball, astronaut, and planet are replaced by a rock, an alien, and an asteroid, respectively. In this worked solution step, the speed of the baseball is assumed to be 0 meters per second at the top of its projectile motion event. Analogical model formulation uses the candidate inferences resulting from this worked solution step to make the modeling decision that the rock's speed at the top of its projectile motion event is 0 meters per second. While each of the worked solution steps results in candidate inferences, analogical model formulation applies only those representing modeling decisions, such as instantiating equations, assuming values, and checking boundary conditions, to the problem. After creating the scenario model, the Companion uses a rule-based problem solver to solve for the sought quantity and select the appropriate multiple choice answer.

The following transfer learning evaluation was designed by the Educational Testing Service and Cycorp to evaluate a Companion's ability to transfer the knowledge necessary to solve AP physics-style problems using analogical model formulation. Given a source task of 20 problems and worked solutions, target learning for six transfer

```

(stepType Step3 DeterminingValueFromContext)
(stepUses Step3 (isa Throwing1 ThrowingAnObject))
(stepUses Step3 (occursNear Throwing1 Ground1))
(stepUses Step3
  (no-GenQuantRelnFrom
    in-ImmersedFully Planet1 Atmosphere))
(stepUses Step3 (objectMovingUpward1 BaseBall1))
...
(stepUses Step3 (direction Upward1 UpDirectly))
(solutionStepResult Step3
  (valueOf
    (AtFn ((QPQuantityFn Speed) BaseBall1)
      (EndFn Upward1)))
  (MetersPerSecond 0)))
...
(groundOf Asteroid1 Ground2)
(performedBy Throwing2 Alien1)
(no-Gen QuantRelnFrom
  in-ImmersedFully Asteroid1 Atmosphere)
(eventOccursNear Throwing2 Ground2)
(objectThrown Throwing2 Rock1)
(querySentenceOfQuery Query2
  (valueOf (AccGravityFn Asteroid1) Acc1))
...

```

Figure 3. A Portion of the Representation of a Source Worked Solution (left) and Corresponding Restyling Problem (right).

(Simplified for readability.)

distances was performed independently as follows. The Companion was given a series of five training sets each consisting of a sequence of four quizzes. Each quiz consists of one problem from each of the four types. After each quiz, the Companion received the worked solutions for the problems on that quiz. After each training set, the Companion's memory was reset. In the transfer condition, the Companion began with the worked solutions to the source problems and worked solutions in its memory. In the nontransfer condition, the Companion began with zero worked solutions.

The learning curves grouped by transfer distance are shown in figure 4. Averaged across all the transfer distances, the Companion achieved a 95.8 percent initial advantage due to the source problems and their worked solutions. On parameterization, extrapolation, restructuring, extending, and restyling problems, the Companion exhibited perfect transfer. That is, the system performed at ceiling (100 percent) given just the source set worked solutions. On composing problems, the system recorded an initial advantage of 75 percent. All of the initial advantages are statistically significant ($p < .01$).

These results illustrate that analogical model formulation enables a system to transfer the model formulation knowledge necessary to solve AP physics-style problems for the six transfer distances shown. The only failures of transfer involved limitations in the verification of analogical inferences. In particular, our system prevented a necessary modeling decision from being made on a set of dif-

ficult problems (25 percent of the composing problems).

Other Evaluations of CBR as Transfer Learning Methods

Other researchers have explored the evaluation of CBR as a transfer learning method. Also using SME for analogical reasoning, Hinrichs and Forbus (2011) transfer learned city management decisions in a turn-based strategy game. The transfer distance in this work involves changing the configuration of tiles around the cities. As another example, CARL (Sharma et al. 2007) integrates CBR with reinforcement learning to play a real-time strategy game. CARL uses reinforcement learning for credit assignment and CBR to estimate a value function for a given task. After learning in the source domain, the case base is used directly in the target. CARL was evaluated on two transfer distances: (1) the target scenario swapped the starting locations of friendly and opposition forces, and (2) the number of scenario entities was at least 50 percent greater in the target scenario than the source. Using a weighted combination of state features as a performance metric, CARL demonstrated a significant positive initial advantage on both transfer distances and an asymptotic advantage on the second transfer distance.

Discussion

As transfer learning methods, CBR approaches have predominately been applied to transfer distances where the representational vocabulary of

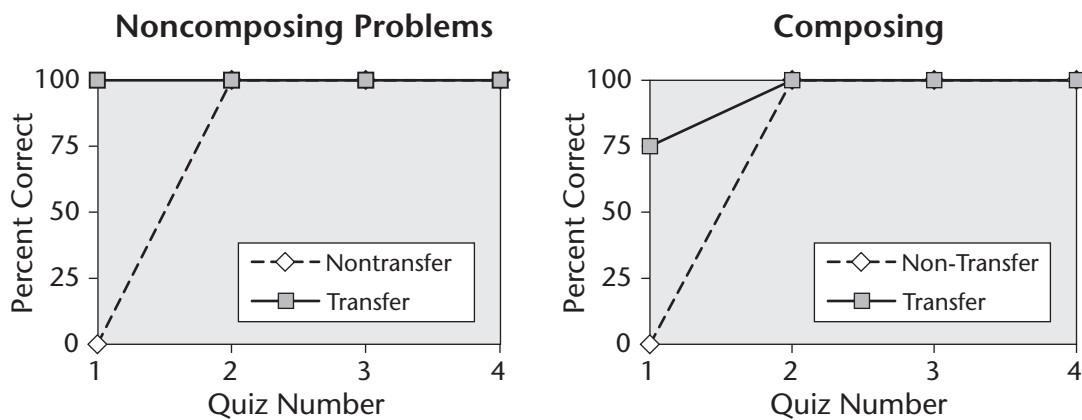


Figure 4. Average Problem-Solving Performance Using Analogical Model Formulation to Solve AP Physics-Style Problems.

the problem (that is, game states and physics scenarios) and what gets reused (that is, possible actions and modeling decisions) are shared across source and target. This permits cases from either source or target to be directly applied to target problems. Most CBR retrieval systems rely on feature similarity. Consequently, identifying and retrieving analogous source instances is very difficult when the vocabulary of the problems and solutions significantly differs between source and target. Consequently, performing transfer learning in a CBR cycle is unlikely to work for such transfer distances.

One of the original motivations for CBR research was to provide greater flexibility than rule-based expert systems for solving unanticipated problems. While measuring this flexibility is difficult, transfer learning provides one approach for empirically evaluating a CBR system by categorizing differences in source and target problems by transfer distance. Transfer learning contributes two important metrics for evaluating the flexibility of CBR systems. First, the initial advantage metric empirically measures the flexibility of the CBR system's retrieval and reuse mechanisms. That is, the transfer distance indicates how similar problems have to be in order to retrieve and reuse a solution. Second, when the system is unable to solve a target problem with the source problems, the learning rate measures the retrieval mechanism's ability to avoid source cases or the CBR system's ability to perform case-based maintenance (Leake et al. 2001). Empirically evaluating the same CBR system across a range of transfer distances, as in the AP physics evaluation above, provides information regarding the flexibility of the CBR system.

CBR for Problem Learning: Rush Football

CBR systems learn by revising and retaining cases as they gain experience. Approaches using CBR for the problem learning transfer learning step maintain a case base for a set of problems (source or target). Here, we present an example of this approach for both source and target problems with separate case bases. The source task, intent recognition (Sukthankar 2007), is to identify the opposing team's intent (that is, play), and the target task is to control a player in the Rush 2008 American football simulation (Rush 2008).⁴ Rush simulates a simplified version of American football with only eight players on each team and whose field is 100 by 63 yards. Figure 5 displays an annotated screenshot from the simulator. The offensive team's objective is to advance the ball into the defensive team's end zone, while the defensive team's objective is to prevent this. The game is divided into downs. Before each down, each team secretly selects a play, which is a set of instructions for each player on the team. Each down ends when the offensive player is tackled by an opposing player or a forward pass is incomplete.

Case-Based Q-Lambda with Intent Recognition

Case-based q-lambda with intent recognition, or CBQL-IR (Aha, Molineaux, and Sukthankar 2009), applies CBR separately to the source task, intent recognition, and the target task, controlling the quarterback. CBQL-IR uses CBR for both the source and target problems. For intent recognition prob-

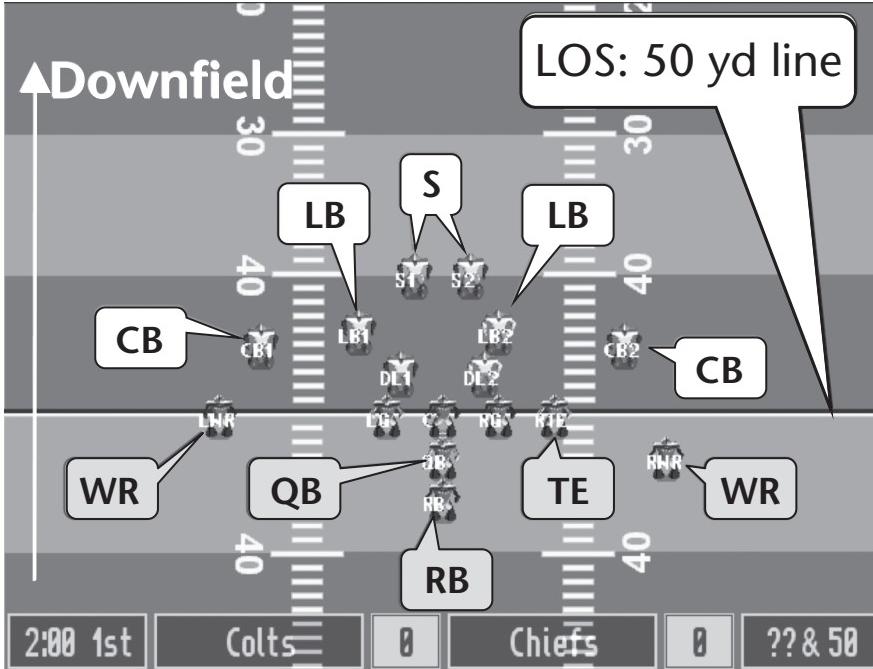


Figure 5. Screenshot at the Beginning of a Down.

The QB is the quarterback and begins the play with ball. The wide receivers (WRs) and tight end (TE) move downfield to receive a pass from the QB. The running back (RB) either takes a handoff from the QB or runs down field to receive a pass.

lems, a k -nearest neighbor, or k -NN (Dasarathy 1991), is used to infer the opponent's play. For the QB control problems, CBR is used to approximate the standard reinforcement learning Q-function.

To perform the source task, intent recognition, CBQL-IR uses a k -nearest neighbor classifier with one example from each of the eight possible defensive plays. Each example includes a feature vector representing the movements of each of the eight defensive players over the first three time steps of the play. As a result of this source task learning, CBQL-IR accurately classifies the current defensive play after three time steps.

The target task is to control the quarterback (QB) by selecting actions, shown in figure 6, at each time step. There are nine actions available to the quarterback: passing the ball to WR1, WR2, TE, or RB, moving in one of four directions, and remaining stationary. The other players all execute the actions described in the offensive and defensive plays. CBQL-IR models this as an RL task with rewards based on the outcome of each play. At each state, CBQL-IR uses the cases with the most similar states to estimate the expected reward for each action. States are represented using two features: the predicted opposing play and the current

time step. At the end of the down, the received reward is used to store new cases, if they are sufficiently different from previous ones, or update the existing similar cases accordingly.

By separating source and target tasks in this manner, the transfer learning evaluation acts as an ablation study, measuring the effects of intent recognition on CBQL-IR's ability to control the QB. In the transfer condition, the predicted defensive play feature was determined by the k -NN classifier learned from training on source problems. In the nontransfer condition, this feature was assigned a random value because it had no knowledge from the source problems.

The results are shown in figure 7. While CBQL-IR begins with identical performance in the two conditions, it learns faster and achieves a higher ceiling in the transfer condition. There is no initial advantage because CBQL-IR acts randomly when there are no stored cases, as is true for any reinforcement learner before training occurs. As detailed in Aha, Molineaux, and Sukthankar (2009), intent recognition can significantly improve the system's task performance for this application. To summarize, after target training, the offense gained on average around 10 yards in

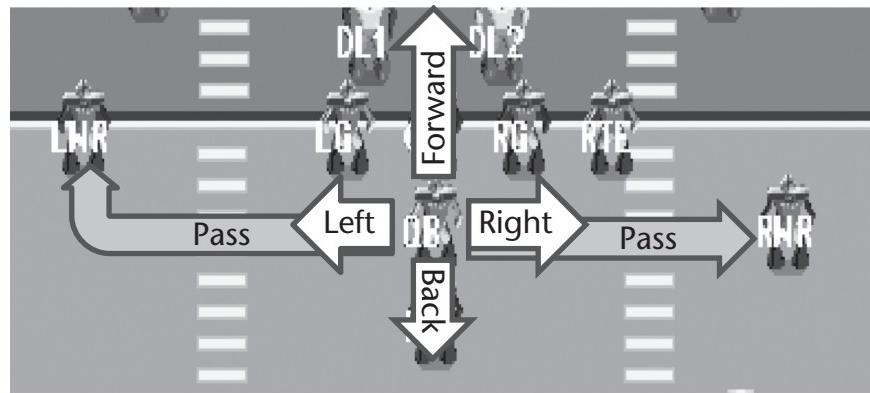


Figure 6. Subset of Actions Available to the QB at Each Time Step.

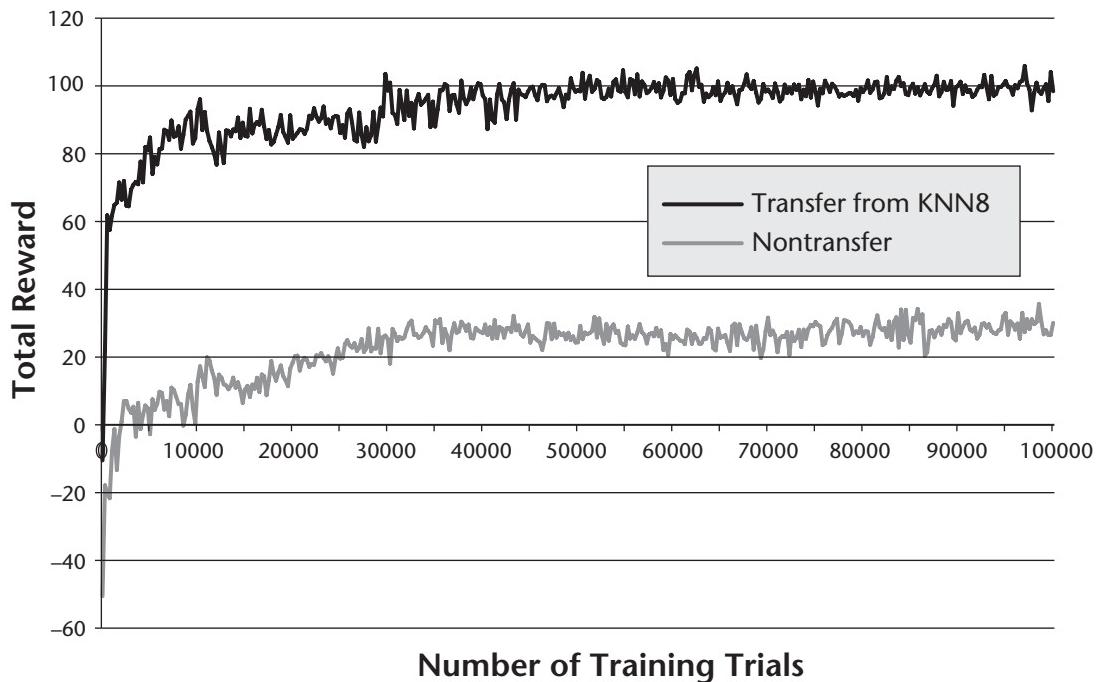


Figure 7. CBQL-IR's Learning Curves for Controlling the Rush 2008 QB on Repeated Plays.

the transfer condition, but only about 3 yards in the nontransfer condition. Therefore, transfer learning measures the effects of integrating intent recognition with controlling an offensive player.

The next section describes other systems that use CBR solely for problem learning within the transfer learning framework.

Other Applications of CBR for Problem Learning

Related research on instance transfer (Pan and Yang 2010) includes approaches that integrate auxiliary (source) data into k-NN classifiers (Wu and Dietterich 2004; Shi et al. 2009). These approaches assume an abundance of source instances and few target instances. When the known target instances are insufficient for classification, these approaches use similar source instances for classification. These approaches are applications of CBR to target task learning because they distinguish source instances from target instances.

Wu and Dietterich (2004) demonstrate the utility of using source instances when the target instances are insufficient. They minimized the weighted sum for two loss functions: one over the target instances, and the other using only the source instances. This results in a classifier in which target instances are weighted higher than source instances for k-NN classification. By including the source instances, they significantly increased accuracies for classifying leaf species. While Wu and Dietterich's approach weights all source cases equally, COITL (Shi et al. 2009) assigns weights to each source instance based on their similarity to the target task. The process of assigning weights to source instances is also an application of CBR to transfer knowledge, and the specifics are discussed later. COITL demonstrates substantial improvements in generalization and robustness over previous approaches on a number of standard datasets. In contrast to using CBR as a transfer learning method, each of these systems maintains a distinction between source and target instances.

Discussion

CBQL-IR illustrates how case-based reasoning can be used in different source and target tasks. Classification of the opponent's play enabled faster learning and higher asymptotic performance when controlling the Rush 2008 quarterback's actions. While the source learning, play classification, informs the problem representation in the target task, player control, the case-based reasoning in the target task only reuses solutions from target cases. This is due to the differences between source and target case representations and solutions. One cannot adapt a play classification directly into the selection of an action.

When the source and target tasks are drastically

different, as in CBQL-IR, then measuring the transfer distance is not applicable. Instead, transfer learning functions as an ablation experiment evaluating the contribution of an individual component in an integrated systems research. However, when the source and target tasks are the same, as in the instance transfer approaches, the transfer distance is noteworthy. In such approaches, the features used for classification may differ among source and target problems, but they need to be derivable from each target instance. Furthermore, the labels must be shared among the source and target.

CBR for Transferring Knowledge: Linear and Rotational Kinematics

Approaches using CBR to transfer knowledge focus on identifying how the source and target relate to each other. To illustrate this approach, we present an application of domain transfer via analogy to learn a mapping between linear and rotational mechanics (Klenk and Forbus 2009b). In this work, an analogical mapping is created between source and target, which supports the transfer of abstract source knowledge. As in the AP physics work, the source and target consists of physics problems and worked solutions. Following the same conventions as the work in the CBR as a Transfer Learning Method section, the problems and worked solutions are represented in predicate calculus using the ResearchCyc ontology. Unlike the AP physics work, the system uses a domain theory of kinematics equations defined using equation schemas, which specify the equation's types, quantities, and applicability conditions. In both the source and target tasks, the system uses the available equation schemas and rule-based problem solving to construct the scenario model and solve for the sought quantity.

Domain Transfer Via Analogy

Domain transfer via analogy (figure 8) learns the equation schemas necessary to solve target problems by cross-domain analogy with a source domain. After failing to solve a target problem, domain transfer via analogy is provided with its worked solution. The inputs to domain transfer via analogy are this worked solution, the source problems' worked solutions, and the source equation schemas. In step 1, domain transfer via analogy uses MAC/FAC with the target worked solution to retrieve an analogous source worked solution, and SME to create a set of correspondences between the types, quantities, and relations of the source and target, called a domain mapping. During this process, SME aligns nonidentical relations using minimal ascension (Falkenhainer 1988), as described in Hinrichs and Forbus (2011).

In step 2, this domain mapping is used to ini-

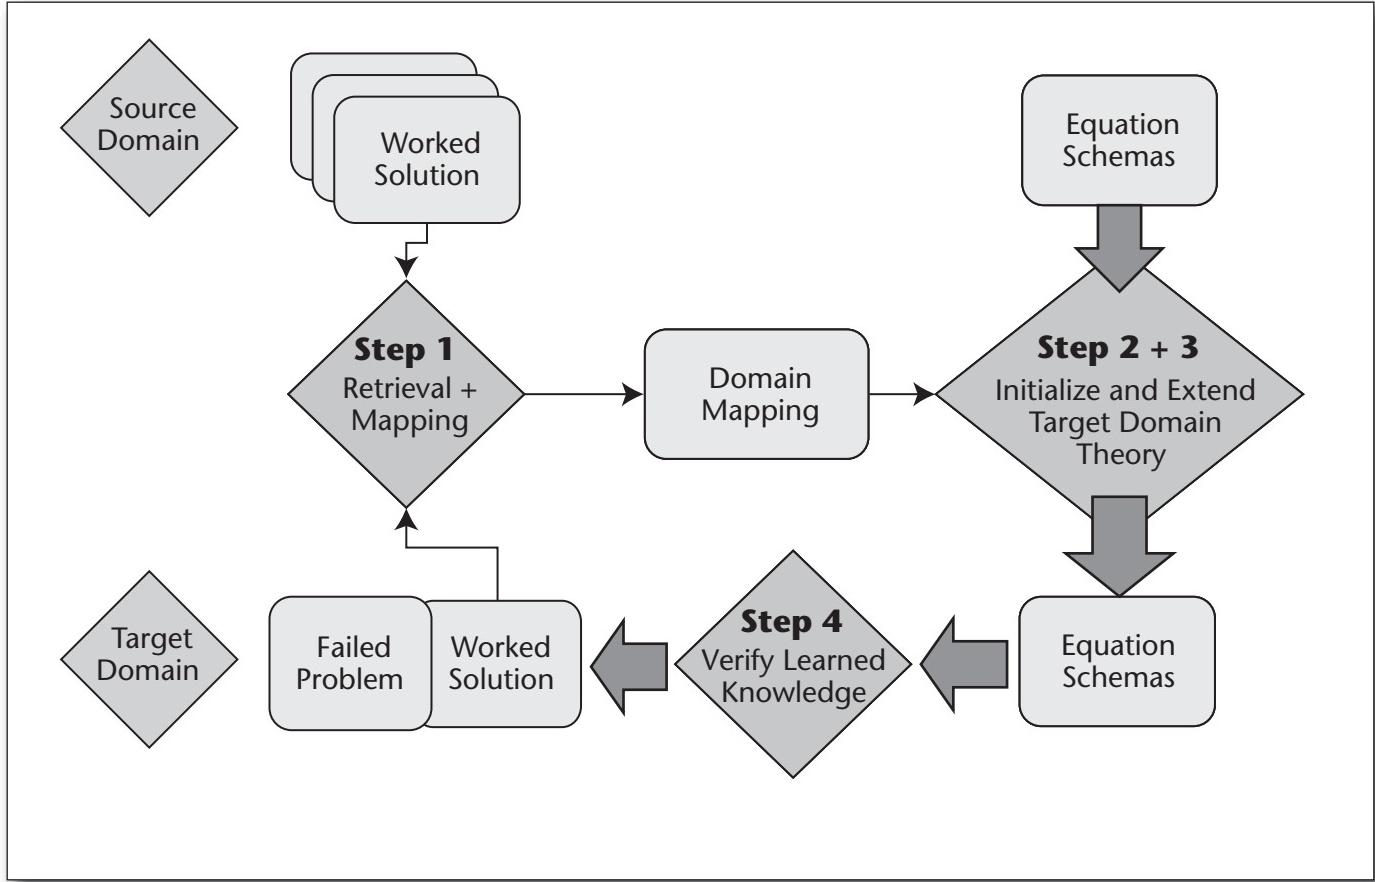


Figure 8. Conceptual Model for Domain Transfer Via Analogy.

tialize the target domain theory. For each source equation schema in the domain mapping, domain transfer via analogy replaces its types, predicates, and quantities with the corresponding expressions from the domain mapping. For example, the linear kinematics (source) domain includes the equation $d = v_i t + .5at^2$, which has an applicability condition stating that the predicate objectTranslating must hold between the object and the movement event of the equation schema. Consider a domain mapping resulting from step 1 that includes a correspondence between the predicates objectTranslating and objectRotating. Therefore, the initialized equation schema's applicability condition is defined in terms of the objectRotating predicate. After all the substitutions, the rotational kinematics (target) domain includes an equation schema defining $\theta = \omega_i t + .5\alpha t^2$. In step 3, the target domain theory is extended through an analogy between the domain theories themselves. This permits domain transfer via analogy to transfer equation schemas not mentioned in the target worked solutions. Finally, because analogical learning is not sound, domain transfer via analogy verifies the transferred equation schemas by attempting to

solve the target problem. If successful, the target equation schemas are assumed to be correct and retained for use in solving future target problems. Otherwise, the transferred knowledge and domain mapping are discarded.

The following transfer learning evaluation was performed to measure the effectiveness of domain transfer via analogy for transferring knowledge between linear and rotational kinematics. Each domain includes five problems and worked solutions. Target learning was measured by running 120 trials representing every possible ordering of the test materials. In the transfer condition, the domain transfer via analogy was provided with the five worked solutions and the equation schemas from the source domain. In the transfer condition, after the system failed to answer a target problem correctly, domain transfer via analogy was invoked with its worked solution. In the nontransfer condition, after failing to solve a target problem, the equation schemas necessary to solve that problem were added directly to the system's target domain theory. Two experiments were performed: one with linear kinematics as the source and rotational kinematics as the target, and vice versa.

The results, shown in figures 9 and 10, graph the percentage of problems the system answered correctly against a problem's position within the trial. In learning rotational mechanics, while domain transfer via analogy transferred the equation schemas necessary to answer correctly the remaining questions after only one problem, the non-transfer condition only answered the last problem on the quiz correctly 80 percent of the time. In learning linear kinematics, after any three problems both the nontransfer and transfer conditions were able to correctly answer the rest of the problems. An analysis of the few linear kinematics transfer failures indicates that increasing the number of subevents and time intervals in the problem representation increases the difficulty in generating a useful domain mapping. The results from both experiments demonstrate that domain transfer via analogy outperformed the nontransfer condition by learning faster and achieved the same or higher ceiling. This work has been expanded to demonstrate transfer between linear mechanical, rotational, electrical, and thermal systems (Klenk and Forbus 2009c).

Other Applications of CBR for Transferring Knowledge

We distinguish these approaches according to whether they perform analogical or instance transfer. Analogical transfer approaches, including domain transfer via analogy, create mappings between the source and target supporting the transfer of abstract source knowledge. Instance transfer approaches identify source instances that are directly applicable to the target task.

Analogical Transfer Methods. As discussed earlier, structure mapping theory (Gentner 1983) defines analogy as an alignment process between two structured representations. The resulting alignment is a mapping between entities and expressions of the two representations, which satisfy the analogical process model's structural, semantic, and pragmatic constraints (Holyoak and Thagard 1989). Different process models (SME, Falkenhainer, Forbus, and Gentner [1989]; LISA, Hummel and Holyoak [2003]; AMBR, Kokinov and Petrov [2001]) employ slightly different versions of these constraints. Analogy research seeks to understand the principles underlying the retrieval of analogous cases, the mapping between them, the inferences suggested, and the integration of analogical reasoning in cognition. This alignment process is essential for retrieval and reuse in CBR. Therefore, we consider the transfer learning methods below as using partial CBR cycles for transfer. These approaches differ with respect to the model of analogy, representations, and transferred knowledge, which we summarize in table 2.

Lui and Stone (2006) present a method for ana-

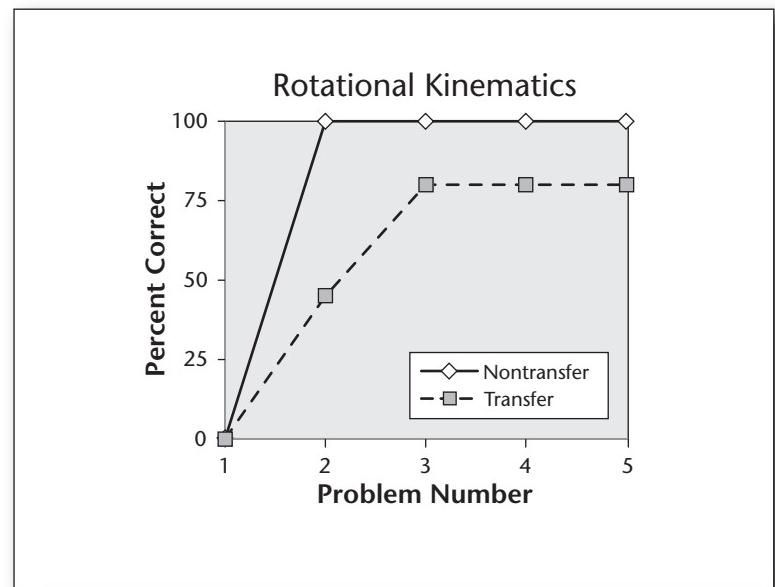


Figure 9. Rotational Kinematics as Target.

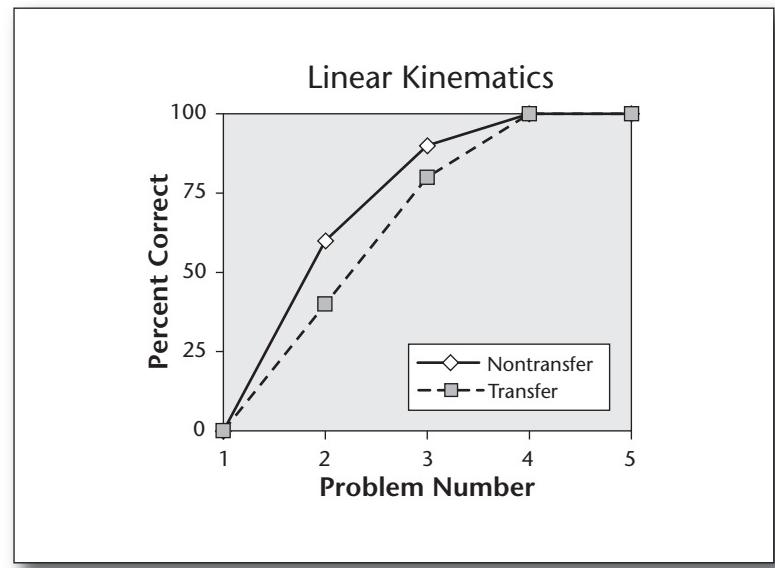


Figure 10. Learning Linear Kinematics as Target.

logical transfer using a version of SME that was optimized to map sets of qualitative dynamic Bayes nets (QDBNs). QDBNs are provided for the source and target RL problems describing the effects of an action. The resulting mapping is used to transfer Q-values between reinforcement learning problems. Transfer over QDBNs improved performance in learning policies for keepaway soccer games of various sizes.

Isomorphism between two graph representations underlies most analogy research. Kuhlmann

	Model of Analogy	Base and Target Representations	Transferred Knowledge
Lui and Stone (2006)	SME-QDBN	QDBNs of Task Models	RL Value Function
Kuhlmann and Stone (2007)	Graph Isomorphism	Canonical GGP Game Rules	RL Value Function
Klenk and Forbus (2009)	SME	Physics Worked Solutions	Equation Schemas
König et al. (2009)	GAMA	Analyzed GGP Game Rules	Hierarchical Skills
Hinrichs and Forbus (2011)	SME	Analyzed GGP Game Rules	HTN-Methods

Table 2. Analogical Transfer Methods for Transfer Learning.

and Stone (2007) transform general game playing, or GGP (Genesereth, Love, and Pell 2005), game rules into a conical form stored as a graph for each game. Next, they find graph isomorphisms between games to reuse Q-value functions learned from playing the source game. Transformations between source and target include different board sizes (3x3 minichess to 4x4 minichess) or inverted goals (winning tac-tac-toe to losing tic-tac-toe).

Hinrichs and Forbus (2011) use the static analysis of GGP game rules to classify game predicates into broad classes including types, quantities, spatial coordinates, and ordinal relations. This information enables the alignment of nonidentical predicates through minimal ascension and metamappings. Using the Companion cognitive architecture, they transfer hierarchical task network methods across five different transfer distances, enabling faster learning in the target problems.

The Icarus cognitive architecture employs goal-driven analogical mapping (GAMA) to translate acquired skills from source problems to related target problems (König et al. 2009). Icarus explains a provided trace from a source problem using concepts from its source domain theory. GAMA uses this trace pragmatically to constrain the mapping between source and target domain theories defined by GGP rules. This mapping enables the transfer of hierarchical skills learned in the source domain, thereby accelerating search in the target domain.

Transfer in each GGP transfer problem requires mapping nonidentical predicates. Each of these approaches includes higher-order relations to constrain the matching process. While the tier-identicality constraint of structure mapping theory (Gentner 1983) for analogical mapping states a strong preference for identical predicates, it also permits nonidentical predicates to match, such as those used in the works described in this section. Future work on analogy from both AI and psychological perspectives will investigate when and how these nonidentical matches may arise.

Instance Transfer Methods. An alternative approach is to use CBR to learn weights for source

instances for use in solving target problems. For example, COITL (Shi et al. 2009) assigns weights to source instances based on their similarity to known target instances. A classifier consisting of known target instances is used to classify each source instance. If the classification succeeds, then the source instance is given the weight of the labeling confidence of the classifier. If the classifier fails, then the source instance's weight is set to zero. Transfer is performed by adding the positive weighted source instances to the target (k-NN) classifier.

Discussion

Domain transfer via analogy creates an analogical mapping between the source and target domains to support the transfer of abstract source knowledge. This enables the system to learn rotational kinematics by analogy with linear kinematics, and vice versa. Using SME and minimal ascension, domain transfer via analogy successfully transfers the kinematics equation schemas necessary for problem solving and demonstrates improved learning over the nontransfer condition. The transfer distance between these domains includes different relations, quantities, and types.

The two types of CBR for knowledge transfer, analogical and instance transfer, are applicable over different transfer distances. As indicated in the previous discussion subsection, instance transfer methods are applicable for transfer distances where source and target feature spaces differ but use the same class labels. Analogical transfer systems are applicable for domains that include different relations in both the problem and solution representations. In fact, the GGP representations for source and target included only a small number of shared relations. Each of the analogical approaches employs higher-order relations to constrain the matching process.

Conclusions

CBR research is the study of how and under what circumstances knowledge from specific cases may be applied to new problems. Therefore, it is implic-

itly tied to transfer learning, which measures the performance difference resulting from a system's experience in a source task on a related target task. We identified three distinct approaches for using case-based reasoning in transfer learning: CBR as a transfer learning method, CBR for problem learning, and CBR for transferring knowledge from source to target domains. Our analysis of these approaches illustrates how transfer learning and case-base reasoning inform one another.

CBR as a transfer learning method: These CBR systems perform transfer learning by accumulating cases and applying them to new problems without any distinction between source and target problems. Currently, these systems have only been applied to the transfer distances where the representation vocabulary for problems and solutions are shared across source and target. In addition to performing near transfer, the transfer learning framework provides these systems with two ways for evaluating their utility. First, by measuring performance as target cases are accumulated, it measures the (potentially negative) impact of the source cases on retrieval and reuse. That is, as the agent accumulates target cases, its retrieval mechanism should ignore the source cases in favor of the more applicable target cases. Second, the division of source and target problems allows the initial advantage metric to empirically assess the CBR system's retrieval and reuse mechanisms for a particular transfer distance.

CBR for problem learning: These integrated systems distinguish source and target cases. As in CBQL-IR, the source and target problem types may be sequential (and differ). CBQL-IR uses the output from the (classification) source task as input to the (reinforcement learning) target task. Transfer learning measures the ability of the CBR problem learner to leverage source knowledge while performing the target task, similar to an ablation study. Alternatively, the source and target problem types may be the same. Here, the source instances act as auxiliary knowledge sources that complement the limited number of target instances. For these systems, transfer learning measures the effects of these auxiliary cases.

CBR for transferring knowledge: In these approaches, a CBR process transfers knowledge from the source for use by the domain learner in the target. Analogical transfer methods have been used for transfer distances in which different relations, types, and quantities define the source and target problems and solutions. These approaches have been integrated with a variety of domain learners for classification, reinforcement learning, hierarchical skill learning, and hierarchical task network method learning. Furthermore, they provide a link between machine learning and cognitive science. We hope this communication continues, as an

interdisciplinary approach to studying intelligence is essential to scientific progress in AI.

In the Transfer Learning subsection, we discussed an important contribution of transfer learning: the categorization of differences between source and target problems by transfer distance. We expect future researchers to define and use transfer distances to better understand the strengths and weaknesses of their CBR approaches. Pan and Yang's (2010) categorizations are a strong step forward for characterizing transfer learning approaches, and we suggest four additional dimensions.

The first is direct applicability of source knowledge. Knowledge learned on source problems is directly applicable if it assists with solving target problems with no additional transformations. For CBQL-IR, the source and target problems differ, but the knowledge gained in the source was from the same frame of reference, and thus was directly applicable. On the other hand, the analogical transfer methods transform the source knowledge for application in the target.

The second dimension is whether the learning problem types are consistent. For example, Wu and Dietterich's (2004) learning task is classification for both source and target problems. As an example of inconsistent types, CBQL-IR's (Aha, Molineaux, and Sukthankar 2009) source task is classification and target task is reinforcement learning. This can impact ease of transferring learned source knowledge.

The third dimension concerns problem labeling. Transfer learning approaches typically assume that target problems will be identified distinctly from source problems. However, some research in CBR as a transfer learning method relaxes the assumption that these different kinds of problems are distinctly identified; all new problems are presented as the same, regardless of whether they are source or target problems.

The fourth dimension concerns multiple representations. In complex domains, such as strategy simulations, AI systems must integrate representations for planning, execution monitoring, and spatial reasoning. The incorporation of multiple representations, such as hierarchical task network methods and spatial diagrams, may assist transfer learning approaches in finding and exploiting similarities between the domains (Klenk 2009). While transfer learning has not explicitly explored this dimension, Klenk et al. (2005) applies analogical model formulation to solving physical reasoning problems involving spatial and conceptual representations.

In summary, research on CBR has varied among these dimensions. Future work remains on measuring how these dimensions characterize transfer distance, and how transfer distance can be used to

predict a given CBR method's transfer learning performance. Transfer learning is not an end of itself; it is an important stepping stone to developing agents that exist over a long period of time and learn to solve a wide variety of problems. While transfer learning methods enable the reuse of knowledge learned in the source to improve future performance, in autonomous situations, there is an additional problem of selecting an analogous source. This problem has received infrequent attention in the transfer learning literature.

In addition to purely autonomous agents, transfer learning is an important technology for collaborative agents that work with users on a wide range of problems. These agents could leverage transfer learning methods directly to exploit user-provided cross-domain analogies, such as "rotational motion is analogous to linear motion" or "a heat pump is like a water pump." Further dialogue between agents and users could explore the analogy between the domains to highlight important correspondences and identify the causes of negative transfer. While difficult to construct, advances in transfer learning may substantially increase the competence of intelligent agents.

Acknowledgements

The authors would like to thank the reviewers for their helpful comments. Matthew Klenk is supported by an NRC Postdoctoral Fellowship.

Notes

1. Transfer learning can also support CBR methods in a variety of ways, but we leave that analysis for future work.
2. For more information, see the CBR Wiki, cbrwiki.fdi.ucm.es/wiki/index.php/Main_Page.
3. CBR for problem learning and CBR to transfer knowledge are not mutually exclusive. An integrated system could use two separate CBR cycles, one for problem learning and one to transfer knowledge, to perform transfer learning.
4. See www.knexusresearch.com/projects/rush.

References

- Aamodt, A., and Plaza, E. 1994. Case-Based Reasoning: Foundational Issues, Methodological Variations, and System Approaches. *AI Communications* 7(1): 39–59.
- Aha, D. W.; Marling, C.; and Watson, I. 2005. Case-Based Reasoning Commentaries: Introduction. *Knowledge Engineering Review* 20(3): 201–202.
- Aha, D. W.; Molineaux, M.; and Sukthankar, G. 2009. Case-Based Reasoning for Transfer Learning. In *Proceedings of the Eighth International Conference on Case-Based Reasoning*, Lecture Notes in Computer Science 5650, 29–44. Berlin: Springer.
- Bransford, J. D.; Brown, A. L.; and Cocking, R. R., eds. 2000. *How People Learn: Brain, Mind, Experience, and School*. Washington, DC: National Academy Press.
- Dasarathy, B. V. 1991. *Nearest Neighbor Norms: NN Pattern Classification Techniques*. Los Alamitos, CA: IEEE Computer Society.
- Falkenhainer, B. 1988. Learning from Physical Analogies, Technical Report UIUCDCS-R-88-1479. University of Illinois, Department of Computer Science, Urbana-Champaign, IL.
- Falkenhainer, B., and Forbus, K. 1991. Compositional Modeling: Finding the Right Model for the Job. *Artificial Intelligence* 51(1–3): 95–143.
- Falkenhainer, B.; Forbus, K. D.; and Gentner, D. 1989. The Structure-Mapping Engine: Algorithm and Examples. *Artificial Intelligence* 41(1): 1–63.
- Forbus, K.; Gentner, D.; and Law, K. 1995. MAC/FAC: A Model of Similarity-Based Retrieval. *Cognitive Science* 19(2): 141–205.
- Forbus, K.; Klenk, M.; and Hinrichs, T. 2009. Companion Cognition Systems: Design Goals and Some Lessons Learned. *IEEE Intelligent Systems* 24(4): 36–46.
- Genesereth, M. R.; Love, N.; and Pell, B. 2005. General Game Playing Overview of the AAAI Competition. *AI Magazine* 26(2): 62–72.
- Gentner, D. 1983. Structure-Mapping: A Theoretical Framework for Analogy. *Cognitive Science* 7(1): 155–170.
- Hinrichs, T. R., and Forbus, K. D. 2007. Analogical Learning in a Turn-Based Strategy Game. In *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence*, 853–858. Menlo Park, CA: AAAI Press.
- Hinrichs, T. R., and Forbus, K. D. 2011. Transfer Learning through Analogy in Games. *AI Magazine* 32(1).
- Holyoak, K. J., and Thagard, P. 1989. A Computational Model of Analogical Problem Solving. In *Similarity and Analogical Reasoning*, ed. S. Vosniadou and A. Ortony. New York: Cambridge University Press.
- Hummel, J. E., and Holyoak, K. J. 2003. A Symbolic-Connectionist Theory of Relational Inference and Generalization. *Psychological Review* 110(2): 220–264.
- Kaelbling, L.; Littman, M.; and Moore, A. 1996. Reinforcement Learning: A Survey. *Journal of Artificial Intelligence Research* 4: 237–285.
- Klenk, M. 2009. Transfer as a Benchmark for Multi-Representational Architectures. In *Multi-Representational Architectures for Human-Level Intelligence: Papers from the AAAI Fall Symposium*. Technical Report FS-09-05. Menlo Park, CA: AAAI Press.
- Klenk, M., and Forbus, K. 2009a. Analogical Model Formulation for Transfer Learning in AP Physics. *Artificial Intelligence* 173(18): 1615–1638.
- Klenk, M., and Forbus, K. 2009b. Domain Transfer Via Cross-Domain Analogy. *Cognitive Systems Research* 10(3): 240–250.
- Klenk, M., and Forbus, K. 2009c. Persistent Mappings in Cross-Domain Analogical Learning of Physics Domains. Paper presented at the 2nd International Analogy Conference, Sofia, Bulgaria, 24–27 July.
- Klenk, M.; Forbus, K.; Tomai, E.; Kim, H.; and Kyckelhahn, B. 2005. Solving Everyday Physical Reasoning Problems by Analogy Using Sketches. In *Proceedings of the 20th National Conference on Artificial Intelligence*. Menlo Park, CA: AAAI Press.
- Kokinov, B., and Petrov, A. 2001. Integrating Memory and Reasoning in Analogy-Making: The AMBR Model. In

- The Analogical Mind: Perspectives from Cognitive Science*, ed. D. Gentner, K. Holyoak, and B. Kokinov. Cambridge, MA: MIT Press.
- Kolodner, J. 1993. *Case-Based Reasoning*. San Mateo, CA: Morgan Kaufmann.
- König, T.; O'Rorke, P.; Shapiro, D.; Choi, D.; Nejati, N.; and Langley, P. 2009. Skill Transfer through Goal-Driven Representation Mapping. *Cognitive Systems Research* 10(3): 270–285.
- Kuhlmann, G., and Stone, P. 2007. Graph-Based Domain Mapping for Transfer Learning in General Games. In *Proceedings of 18th European Conference on Machine Learning*, Lecture Notes in Computer Science. Berlin: Springer.
- Leake, D.; Smyth, B.; Wilson, D.; and Yang, Q. 2001. Introduction to the Special Issue on Maintaining Case-Based Reasoning Systems. *Computational Intelligence* 17(2): 193–195.
- López de Mántaras, R.; McSherry, D.; Bridge, D. G.; Leake, D. B.; Smyth, B.; Craw, S.; Faltings, B.; Maher, M. L.; Cox, M. T.; Forbus, K. D.; Keane, M.; Aamodt, A.; and Watson, I. D. 2005. Retrieval, Reuse, Revision and Retention in Case-Based Reasoning. *Knowledge Engineering Review* 20(3): 215–240.
- Liu, Y., and Stone, P. 2006. Value-Function-Based Transfer for Reinforcement Learning Using Structure Mapping. In *Proceedings of the Twenty-First National Conference on Artificial Intelligence*, 415–20. Menlo Park, CA: AAAI Press.
- Nejati, N.; Langley, P.; and König, T. 2006. Learning Hierarchical Task Networks by Observation. In *Proceedings of the 23rd International Conference on Machine Learning*, 665–672. New York: Association for Computing Machinery.
- Pan, S., and Yang, Q. 2010. A Survey on Transfer Learning. *IEEE Transactions on Knowledge and Data Engineering*. 22(10): 1345–1359
- Perkins, D. N., and Salomon, G. 1994. Transfer of Learning. In *International Handbook of Educational Research*, ed. T. Husen and T. N. Postelwhite, 6452–6457. Oxford, UK: Pergamon Press.
- Rickel, J., and Porter, B. 1994. Automated Modeling for Answering Prediction Questions: Selecting the Time Scale and Boundary. In *Proceedings of the 12th National Conference on Artificial Intelligence*. Menlo Park, CA: AAAI Press.
- Schank, R. 1982. *Dynamic Memory: A Theory of Learning in Computers and People*. New York: Cambridge University Press.
- Sharma, M.; Holmes, M.; Santamaría, J. C.; Irani, A.; Isbell Jr., C. L.; and Ram, A. 2007. Transfer Learning in Real-Time Strategy Games Using Hybrid CBR/RL. In *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence*, 1041–1046. Menlo Park, CA: AAAI Press.
- Shi, Y.; Lan, Z.; Liu, W.; and Bi, W. 2009. Extending Semi-Supervised Learning Methods for Inductive Transfer Learning. In *Proceedings of the 2009 IEEE International Conference on Data Mining (ICDM)*. Piscataway, NJ: Institute of Electrical and Electronics Engineers.
- Sukthankar, G. 2007. Activity Recognition for Agent Teams. Carnegie-Mellon University Technical Report CMU-RI-07-23, Pittsburgh, PA.
- Sutton, R., and Barto, A. 1998. *Introduction to Reinforcement Learning*. Cambridge, MA: The MIT Press.
- Sutton, R.; Precup, D.; and Singh, S. 1999. Between MDPs and Semi-MDPs: A Framework for Temporal Abstraction in Reinforcement Learning. *Artificial Intelligence* 112(1–2): 181–211.
- Taylor, M., and Stone, P. 2009. Transfer Learning for Reinforcement Learning Domains: A Survey. *Journal of Machine Learning Research* 10(1): 1633–1685.
- Thorndike, E. L., and Woodworth, R. S. 1901. The Influence of Improvement in One Mental Function upon the Efficiency of Other Functions (I). *Psychological Review* 8(3): 247–261.
- Thrun, S., and Pratt, L. Y., eds. 1998. *Learning to Learn*. Kluwer Academic Publishers, Boston, MA.
- Torrey, L., and Shavlik, J. 2009. Transfer Learning. In *Handbook of Research on Machine Learning Applications*, ed. E. Soria, J. Martin, R. Magdalena, M. Martinez, and A. Serano. Hershey, PA: IGI Global.
- Watkins, C. J. C. H. 1989. Learning from Delayed Rewards. PhD dissertation, University of Cambridge, England.
- Wu, P., and Dietterich, T. 2004. Improving SVM Accuracy by Training on Auxiliary Data Sources, In *Proceedings of the 21st International Conference on Machine Learning*. New York: Association for Computing Machinery.

Matthew Klenk is an NRC postdoctoral research associate at the Navy Center for Applied Research in Artificial Intelligence at the U.S. Naval Research Laboratory. He received his PhD in electrical engineering and computer science from Northwestern University in 2009. His research interests include analogy, multimodal interaction, spatial reasoning, and integrated planning and execution systems. He is an active member of AAAI and the Cognitive Science Society.

David W. Aha is the supervisor for the Adaptive Systems Section (Code 5514) within the Naval Research Laboratory's Navy Center for Applied Research in Artificial Intelligence. He received his PhD from the University of California, Irvine in 1990. His research interests include case-based reasoning, mixed-initiative reasoning, machine learning, planning, text analysis, and related topics pertaining to intelligent decision aids. He was an AAAI councilor, founded the UCI Repository of Machine Learning Databases, cofounded the AI Video Competition, and has received three best paper awards. Aha has coorganized 21 international research events, coedited three special journal issues on AI topics, participated on 14 dissertation committees, serves on the editorial boards for three journals, and serves annually on the program committees for several conferences, workshops, and doctoral symposiums.

Matthew Molineaux is a senior computer scientist at Knexus Research Corporation. He holds an MS in computer science from DePaul University and is currently pursuing a PhD at George Mason University. He has designed, developed, and integrated a number of adaptive autonomous agents and systems that operate in complex environments. His primary interests are in developing new reasoning techniques for autonomous agents, primarily for automated planning, abductive explanation, online learning, and goal reasoning. He publishes on topics including case-based reasoning, transfer learning, AI and games, goal-directed autonomy, and automated planning.